

NBER WORKING PAPER SERIES

OPTIMAL MITIGATION POLICIES IN A PANDEMIC:
SOCIAL DISTANCING AND WORKING FROM HOME

Callum J. Jones
Thomas Philippon
Venky Venkateswaran

Working Paper 26984
<http://www.nber.org/papers/w26984>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

We are grateful for comments from seminar participants at the University of Chicago. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, its management, or the National Bureau of Economic Research. First version: March 31, 2020.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Callum J. Jones, Thomas Philippon, and Venky Venkateswaran. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home
Callum J. Jones, Thomas Philippon, and Venky Venkateswaran
NBER Working Paper No. 26984
April 2020
JEL No. E2,E6,I1

ABSTRACT

We study the response of an economy to an unexpected epidemic. Households mitigate the spread of the disease by reducing consumption, reducing hours worked, and working from home. Working from home is subject to learning-by-doing and the capacity of the health care system is limited. A social planner worries about two externalities, an infection externality and a healthcare congestion externality. Private agents' mitigation incentives are weak and biased. We show that private safety incentives can even decline at the onset of the epidemic. The planner, on the other hand, implements front-loaded mitigation policies and encourages working from home immediately. In our calibration, assuming a CFR of 1% and an initial infection rate of 0.1%, private mitigation reduces the cumulative death rate from 2.5% of the initially susceptible population to about 1.75%. The planner optimally imposes a drastic suppression policy and reduces the death rate to 0.15% at the cost of an initial drop in consumption of around 25%.

Callum J. Jones
International Monetary Fund
700 19th St NW
Washington, DC 20431
jonescallum@gmail.com

Thomas Philippon
New York University
Stern School of Business
44 West 4th Street, Suite 9-190
New York, NY 10012-1126
and NBER
tphilipp@stern.nyu.edu

Venky Venkateswaran
Stern School of Business
New York University
7-81 44 West 4th Street
New York, NY 10012
and Federal Reserve Bank of Minneapolis
and also NBER
vvenkate@stern.nyu.edu

1 Introduction

The response to the Covid-19 crisis highlights the tension between health and economic outcomes. The containment measures that help slow the spread of the virus are likely to reinforce the economic downturn. Policy makers are therefore faced with difficult decisions.

We propose a simple extension of the neoclassical model to quantify the tradeoffs and guide policy. We are particularly interested in understanding the design of the policy response. When will the private sector engineer the right response, and when is there a need for policy intervention? Which measures should be front-loaded and which ones should ramp up as the contagion progresses?

Our model has two building blocks: one for the dynamics of contagion, and one for consumption and production, including mitigation strategies such as the decision to work from home. Our starting point is the classic *SIR* model of contagion used by public health specialists. Atkeson (2020b) provides a clear summary of this class of model. In a population of initial size N , the epidemiological state is given by the numbers of Susceptible (S), Infected (I), and Recovered (R) people. By definition, the cumulative number of deaths is $D = N - S - I - R$. Infected people transmit the virus to susceptible people at a rate that depends on the nature of the virus and on the frequency of social interactions. Containment, testing, and social distancing reduce this latter factor. The rates of recovery (transitions from I to R), morbidity (I becoming severely or critically sick) and mortality (transition from I to D) depend on the nature of the virus and on the quality of health care services. The quality of health services depends on the capacity of health care providers (ICU beds, ventilators) and the number of sick people. The economic side of the model focuses on three key decisions: consumption, labor supply, and working from home. We use a standard model where members of large households jointly make these decisions. We make three important assumptions: (i) the consumption of (some) goods and services increases the risk of contagion; (ii) going to work also increases the risk of contagion; (iii) working from home involves learning by doing.

We can then study how the private sector reacts to the announcement of an outbreak and how a government should intervene. Upon learning of the risks posed by the virus, households change their labor supply and consumption patterns. They cut spending and labor supply in proportion to the risk of infection, which – all else equal – is proportional to the fraction of infected agents I/N . Households only take into account the risk that they become infected, not the risk that they infect others, therefore their mitigation efforts are lower than what would be socially optimal. This infection

externality is well understood in the epidemiology literature. The other important externality is the congestion externality in the healthcare system. When hospitals are overwhelmed the risk of death increases but agents do not internalize their impact on the risk of others.

We obtain interesting results when we compare the timing of mitigation. The planner wants to front load these efforts compared to the private sector. This is especially true with working from home. Upon learning about the disease, the planner asks people to start working from home so that mitigation is less costly at the peak of the crisis. The period when I/N is low is a good time to learn how to work from home. This lowers the cost of social distancing when we reach the peak of contagion risks.

The contagion and congestion externalities drive a large wedge between private decisions and the socially efficient allocation. If a private agent knows that she is likely to be infected in the future, this reduces her incentives to be careful today. We call this effect the fatalism effect. The planner on the other hand, worries about future infections. When the congestion externality is large, we show that there can even be a front-running incentive. This happens when the private marginal value of being infected becomes larger than the private marginal value of being susceptible. A private agent who thinks that hospitals are likely to be overrun in a few weeks might rationally conclude that it would be better to become sick sooner rather than later.

Literature. Our paper relates to the literature on contagion dynamics (Diekmann and Heesterbeek, 2000). We refer to the reader to Atkeson (2020b) for a recent discussion. Berger et al. (2020) show that testing can reduce the economic cost of mitigation policies as well as reduce the congestion in the health care system. Baker et al. (2020) document the early consumption response of households in the US.

The most closely related papers are Barro et al. (2020), Eichenbaum et al. (2020) and Alvarez et al. (2020). Barro et al. (2020) and Correia et al. (2020) draw lessons from the 1918 flu epidemic. Barro et al. (2020) find a high death rate (about 40 million people, 2% of the population at the time) and a large but not extreme impact on the economy (cumulative loss in GDP per capita of 6% over 3 years). The impact on the stock market was small. Correia et al. (2020) find that early interventions help protect health and economic outcomes.

Our model shares with Eichenbaum et al. (2020) the idea of embedding *SIR* dynamics in a simple DSGE model. The SIR model is the same, but some differences come from the DSGE model.

Eichenbaum et al. (2020) consider hand-to-mouth households while we use a shopper/worker framework à la Lucas and Stokey (1987). Compared to Eichenbaum et al. (2020) we seem to find that optimal interventions are more front-loaded. Apart from this difference, we emphasize the role of learning-by-doing in working from home which adds an important dynamic choice. The planner invests in the new technology to mitigate future disruptions. We also explain the dynamic tension between the planner and the private sector and we describe a fatalism bias in private incentives.

Alvarez et al. (2020) study a lockdown planning problem under *SIR* dynamics. They assume risk neutral agents and a linear lockdown technology. They find that the congestion externality plays an important role in shaping the policy response and that the planner front-loads the effort. Our planner does something similar but takes into account the desire for consumption smoothing.

2 Benchmark Model

2.1 Households

There is a continuum of mass N of households. Each household is of size 1 and the utility of the household is

$$U = \sum_{t=0}^{\infty} \beta^t u(c_t, l_t; i_t, d_t)$$

where c_t is per-capita consumption and l_t is labor supplied by those who are alive and not sick. The household starts with a continuum of mass 1 of family members, all of them susceptible to the disease. At any time $t > 0$ we denote by s_t , i_t and d_t the numbers of susceptible, infected and dead people. The size of the household at time t is therefore $1 - d_t$. If per capita consumption is c_t then household consumption is $(1 - d_t) c_t$. Among the i_t infected members, κi_t are too sick to work. The labor force at time t is therefore $1 - d_t - \kappa i_t$, and household labor supply is $(1 - d_t - \kappa i_t) l_t$. The number of household members who have recovered from the disease is $r_t = 1 - s_t - i_t - d_t$. In the quantitative applications below, below we use the functional form

$$u(c_t, l_t; i_t, d_t) = (1 - d_t - \kappa i_t) \left(\log(c_t) - \frac{l_t^{1+\eta}}{1+\eta} \right) + \kappa i_t (\log(c_t) - u_\kappa) - u_d d_t$$

where u_κ is the disutility from being sick and u_d the disutility from death which includes lost consumption and the psychological cost on surviving members.¹ For simplicity we assume that sickness

¹Formally $u_d = \text{PsyCost} - \log(c_d)$ where c_d is the consumption equivalent in death. Technically we cannot set $c_d = 0$ with log preferences but u_d is a large number.

does not change the marginal utility from consumption therefore c is the same for all alive members of the household. The variables s , i and d evolve according a standard *SIR* model described below. We use a Lucas and Stokey (1987) approach. At the beginning of time t each household decides how much to consume c_t (per capita) and how much each able-bodied member should work l_t . We have normalized the disutility of labor so that $l = c = 1$ before the epidemic starts.

Households understand that they can become infected by shopping and by going to work. We compute infection in two steps. First we define exposure levels for shoppers and for workers. Then we aggregate these into one infection rate at the household level. Finally we take into account the stochastic arrival of a vaccine by adjusting the discount factor β . Formally, we assume an exogenous arrival rate for a cure to the disease. By a cure we mean both a vaccine and a treatment for the currently sick. Under this simplifying assumption the economy jumps back to $l = c = 1$ when a cure is found. We can therefore focus on the stochastic path before a cure is found. Let $\tilde{\beta}$ be the pure time discount rate and ν the likelihood of a vaccine. We define $\beta = \tilde{\beta}(1 - \nu)$ along the no-cure path.

2.2 Shopping

Household members can get infected by shopping. We define consumption (shopping) exposure as

$$e^c c_t C_t,$$

where e^c measures the sensitivity of exposure to consumption and C_t is aggregate consumption, all relative to a steady state value normalized to one. The idea behind this equation is that household members go on shopping trips. We assume that shopping trips scale up with consumption and that, for a given level of aggregate consumption, exposure is proportional to shopping trips. This functional form captures the notion of crowds in shopping mall as well as in public transportation. We study heterogeneity across sectors later in the paper.

2.3 Production

Production uses only labor, but a key feature of our model is the distinction between hours supplied by able bodied workers l_t and effective labor supply \hat{l}_t per household. Effective labor supply is

$$\hat{l}_t = (1 - d_t - \kappa i_t) \left(l_t - \frac{\chi_t}{2} (m_t)^2 \right)$$

The first term captures the fact that the number of valid household member is decreased by death and sickness. The second term captures the cost of implementing mitigation strategies, denoted m_t (e.g., working from home at least some of the time). The benefit of such strategies is a reduction in the risk of infection. Exposure at work for household members working is given by

$$e^l (1 - m_t) l_t (1 - M_t) L_t$$

Working-from-home is subject to learning by doing, so the cost declines with accumulated experience:

$$\chi_t = \chi(\bar{m}_t)$$

where \bar{m}_t is the stock of accumulated mitigation

$$\bar{m}_{t+1} = \bar{m}_t + m_t.$$

The function χ is positive, decreasing, and convex. In the quantitative analysis, we assume the following functional form

$$\chi(\bar{m}_t) = \bar{\chi} (1 - \Delta_\chi (1 - e^{-\bar{m}_t}))$$

The cost initially is $\chi(0) = \bar{\chi} > 0$ and then falls over time as people figure out how to work from home.

The aggregate resource constraint is

$$Y_t = C_t = \hat{L}_t = N \hat{l}_t$$

In our basic model, we ignore the issue of firm heterogeneity and market power. Therefore price is equal to marginal cost

$$P_t = W_t = 1$$

where W_t is the wage per unit of effective labor, which we normalize to one.

2.4 Income and Contagion

At the end of each period, household members regroup and share income and consumption. Household labor income is $W_t \hat{l}_t = \hat{l}_t$ and the budget constraint is

$$(1 - d_t) c_t + \frac{b_{t+1}}{1 + r_t} \leq b_t + \hat{l}_t$$

Household exposure is

$$e_t = \bar{e} + (1 - d_t) e^c c_t C_t + (1 - d_t - \kappa i_t) e^l (1 - m_t) l_t (1 - M_t) L_t,$$

where e_0 is baseline exposure, independent of market activities. Contagion dynamics follow a modified *SIR* model (the benchmark model is explained in the appendix):

$$\begin{aligned} s_{t+1} &= s_t - \gamma e_t \frac{I_t}{N} s_t \\ i_{t+1} &= \gamma e_t \frac{I_t}{N} s_t + (1 - \rho) i_t - \delta_t \kappa i_t \\ d_{t+1} &= d_t + \delta_t \kappa i_t \\ r_{t+1} &= r_t + \rho i_t \end{aligned}$$

where γ is the infection rate per unit of exposure, ρ the recovery rate, κ the probability of being sick conditional on infection, and δ_t the mortality rate of sick patients. In the standard *SIR* model γ is constant. In our model it depends on exposure and therefore on mitigation strategies. The parameter δ_t increases when the health system is overwhelmed.

2.5 Market Clearing and Aggregate Dynamics

Infection dynamics for the the entire population are simply given by the *SIR* system above with aggregate variable $I_t = N i_t$, and so on. The aggregate labor force is $N(1 - \kappa i_t - d_t) l_t$ and total consumption is $N(1 - d_t) c_t$. The market clearing conditions are

$$(1 - d_t) c_t = \hat{l}_t,$$

and for the bond market

$$b_t = 0$$

Finally the mortality rate is described by an increasing function:

$$\delta_t = \delta(\kappa I_t)$$

3 Decentralized equilibrium

3.1 Equilibrium Conditions

Since our model reduces to a representative household model and since $b = 0$ in equilibrium, we simply omit b from the value function. The household's recursive problem is

$$V_t(i_t, s_t, d_t, \bar{m}_t) = \max_{c_t, l_t, m_t} u(c_t, l_t; i_t, d_t) + \beta V_{t+1}(i_{t+1}, d_{t+1}, s_{t+1}, \bar{m}_{t+1})$$

where $\bar{m}_{t+1} = \bar{m}_t + m_t$ and

$$u(c_t, l_t; i_t, d_t) = (1 - d_t) \log(c_t) - (1 - d_t - \kappa i_t) \frac{l_t^{1+\eta}}{1+\eta} - u_\kappa \kappa i_t - u_d d_t$$

where effective labor is

$$\begin{aligned} \hat{l}_t &= (1 - d_t - \kappa i_t) \left(l_t - \frac{\chi(\bar{m}_t)}{2} (m_t)^2 \right) \\ \chi(\bar{m}_t) &= \bar{\chi} (1 - \Delta_\chi (1 - e^{-\bar{m}_t})) \end{aligned}$$

The Lagrangian is then:

$$\begin{aligned} V_t &= u(c_t, l_t; i_t, d_t) + \beta V_{t+1} + \lambda_t \left(\hat{l}_t + b_t - (1 - d_t) c_t - \frac{b_{t+1}}{1 + r_t} \right) \\ &\quad + \lambda_{e,t} \left(e_t - \bar{e} - (1 - d_t) e^c c_t \textcolor{red}{C}_t - (1 - d_t - \kappa i_t) e^l (1 - m_t) l_t (1 - \textcolor{red}{M}_t) \textcolor{red}{L}_t \right) \\ &\quad + \lambda_{i,t} \left(i_{t+1} - \gamma e_t \frac{\textcolor{red}{I}_t}{N} s_t - (1 - \rho) i_t + \textcolor{red}{\delta}_t \kappa i_t \right) \\ &\quad + \lambda_{s,t} \left(s_{t+1} - s_t + \gamma e_t \frac{\textcolor{red}{I}_t}{N} s_t \right) \\ &\quad + \lambda_{d,t} (d_{t+1} - d_t - \textcolor{red}{\delta}_t \kappa i_t) \end{aligned}$$

We highlight in red the externalities, from infection and from congestion. The first order conditions for consumption and labor are then

$$\begin{aligned}
c_t : c_t^{-1} &= \lambda_t + \lambda_{e,t} e^c C_t \\
l_t : l_t^\eta &= \lambda_t - \lambda_{e,t} e^l (1 - m_t) (1 - M_t) L_t \\
m_t : \lambda_t \chi_t m_t &= \frac{\beta V_{\bar{m},t+1}}{1 - d_t - \kappa i_t} + \lambda_{e,t} e^l l_t (1 - M_t) L_t
\end{aligned}$$

The remaining first order conditions are

$$\begin{aligned}
e_t : \lambda_{e,t} &= (\lambda_{i,t} - \lambda_{s,t}) \gamma \frac{I_t}{N} s_t \\
i_{t+1} : \lambda_{i,t} &= -\beta V_{i,t+1} \\
s_{t+1} : \lambda_{s,t} &= -\beta V_{s,t+1} \\
d_{t+1} : \lambda_{d,t} &= -\beta V_{d,t+1}
\end{aligned}$$

The envelope conditions are

$$\begin{aligned}
V_{i,t} &= \kappa \frac{l_t^{1+\eta}}{1+\eta} - \kappa u_\kappa - \kappa \lambda_t \left(l_t - \frac{\chi_t (m_t)^2}{2} \right) + \lambda_{e,t} \kappa e^l (1 - m_t) l_t (1 - M_t) L_t - (1 - \rho) \lambda_{i,t} + \delta_t \kappa (\lambda_{i,t} - \lambda_{d,t}) \\
V_{s,t} &= (\lambda_{s,t} - \lambda_{i,t}) \gamma e_t \frac{I_t}{N} - \lambda_{s,t} \\
V_{d,t} &= \frac{l_t^{1+\eta}}{1+\eta} - \log(c_t) - u_d - \lambda_t \left(l_t - \chi_t (m_t)^2 - c_t \right) + \lambda_{e,t} \left(e^c c_t C_t + e^l (1 - m_t) l_t (1 - M_t) L_t \right) - \lambda_{d,t} \\
V_{\bar{m},t} &= \beta V_{\bar{m},t+1} + \lambda_t (1 - d_t - \kappa i_t) \frac{\bar{\chi}}{2} \Delta_\chi e^{\bar{m}_t} (m_t)^2
\end{aligned}$$

3.2 Equilibrium with Exogenous Infections

To simplify the notation, we normalize $N = 1$, so we should think of our values as being per-capita pre-infection. When there is no risk of contagion – $i_t = 0$, $\lambda_{e,t} = 0$ and $V_{\bar{m},t+1} = 0$ – then we have $m_t = 0$ and from the optimal consumption and labor supply

$$c_t^{-1} = l_t^\eta$$

Since $m = 0$ we have $\hat{l}_t = l_t$ so market clearing is simply

$$c_t = l_t.$$

Combining these two conditions we get

$$c_t = l_t = 1$$

The pre-infection economy is always in steady state.

Consider now an economy with exogenous SIR dynamics: $e^c = e^l = 0$. This implies $m_t = 0$ and

$$c_t^{-1} = l_t^\eta.$$

Market clearing requires

$$(1 - d_t) c_t = (1 - d_t - \kappa i_t) l_t$$

therefore

$$l_t^{1+\eta} = 1 + \frac{\kappa i_t}{1 - d_t - \kappa i_t}$$

The labor supply of valid workers increases to compensate for the reduced productivity of the sick.

Per capita consumption is

$$c_t = \left(\frac{1 - d_t}{1 - d_t - \kappa i_t} \right)^{-\frac{\eta}{1+\eta}}$$

As long as $\eta > 0$ consumption per capita decreases. Aggregate GDP decreases because of lost labor productivity and deaths.

The *SIR* system is independent from the economic equilibrium. As described in the Appendix, the share of infected agents I_t increases, reaches a maximum and converges to 0 in the long run. Assuming a constant δ , the long run solution solves

$$\log \left(\frac{S_\infty}{1 - I_0} \right) = -\frac{\gamma \bar{e}}{\rho + \delta \kappa} \left(\frac{1 - S_\infty}{N} \right),$$

and

$$D_\infty = \frac{\delta \kappa}{\delta \kappa + \rho} (1 - S_\infty).$$

When the congestion externality arises and δ_t increases, then we cannot obtain a closed-form solution for the long run death rate but the qualitative results are unchanged. The following proposition

summarizes our results.

Proposition 1. *When contagion does not depend on economic activity ($e^c = e^l = 0$), the share of infected agents I_t increases, reaches a maximum and converges to 0 in the long run. The long run death rate is given by $D_\infty = \frac{\delta\kappa}{\delta\kappa+\rho}(1 - S_\infty)$ where the long run share of uninfected agents solves $\log\left(\frac{S_\infty}{1-I_0}\right) = -\frac{\gamma\bar{e}}{\rho+\delta\kappa}\left(\frac{1-S_\infty}{N}\right)$. Along the transition path, labor supply of able-bodied workers follows the infection rate while per-capita consumption moves in the opposite direction as $c_t = \left(1 - \frac{\kappa i_t}{1-d_t}\right)^{\frac{\eta}{1+\eta}}$.*

3.3 Private Incentives for Mitigation

The key object that summarizes mitigation incentives is the shadow value of exposure:

$$\lambda_{e,t} = (\lambda_{i,t} - \lambda_{s,t}) \gamma \frac{I_t}{N} s_t = \beta (V_{s,t+1} - V_{i,t+1}) \gamma \frac{I_t}{N} s_t$$

Thus, mitigation incentives are the product of two terms – the marginal value of avoiding infection today $\beta (V_{s,t+1} - V_{i,t+1})$ and the current infection risk from exposure, $\gamma \frac{I_t}{N} s_t$. The current infection risk is straightforward to understand. We use an approximation to analyze the first term. First, note that

$$\begin{aligned} V_{s,t+1} &= \beta V_{i,t+2} \gamma e_{t+1} \frac{I_{t+1}}{N} + \left(1 - \gamma e_{t+1} \frac{I_{t+1}}{N}\right) \beta V_{s,t+2} \\ &= \beta V_{i,t+2} \gamma e_{t+1} \frac{I_{t+1}}{N} + \sum_{s=1}^{\infty} \beta^s \left(\prod_{\tau=0}^{s-1} \left(1 - \gamma e_{t+1+\tau} \frac{I_{t+1+\tau}}{N}\right) \right) \gamma e_{t+1+s} \frac{I_{t+1+s}}{N} V_{i,t+1+s} \end{aligned}$$

i.e. the marginal utility of staying susceptible (relative to a fully recovered) is simply the expected discounted value of being infected in the future. The marginal utility of infection is

$$\begin{aligned} V_{i,t+1} &= -\kappa \left[u_\kappa + \lambda_{t+1} \left(l_{t+1} - \frac{\chi_{t+1} (m_t)^2}{2} \right) - \frac{l_{t+1}^{1+\eta}}{1+\eta} + \lambda_{e,t+1} e^l (1 - m_{t+1}) l_{t+1} (1 - M_{t+1}) L_{t+1} \right] \\ &\quad + \delta_{t+1} \kappa \beta V_{d,t+2} + \beta (1 - \rho - \delta_{t+1} \kappa) V_{i,t+2} \end{aligned} \tag{1}$$

The term in square brackets reflects the effects of being sick in period $t+1$ – the direct costs (u_κ), the net cost of not being able to work and reduction in exposure risk. The $\delta_{t+1} \kappa \beta V_{d,t+2}$ captures the expected cost of mortality in $t+1$. The marginal utility of a dead member (again, relative to a fully

recovered)

$$V_{d,t} = -u_d - \left[\log(c_t) - \frac{l_t^{1+\eta}}{1+\eta} + \lambda_t \left(l_t - \frac{\chi_t}{2} (m_t)^2 - c_t \right) - \lambda_{e,t} \left(e^c c_t C_t + e^l (1 - m_t) l_t (1 - M_t) L_t \right) \right] + \beta V_{d,t+1} \quad (2)$$

Now, if u_d is large (as will turn out to be the case in our calibration), it will dominate the other terms in the marginal utility expressions (1) - (2). In other words, the marginal utility of death is well approximated by:

$$V_{d,t} \approx -\frac{u_d}{1-\beta}$$

and that of infection by,

$$V_{i,t+1} \approx -\frac{\beta u_d \kappa}{1-\beta} \left[\delta_{t+1} + \sum_{s=1}^{\infty} \beta^s \left(\prod_{\tau=0}^{s-1} (1 - \rho - \delta_{t+1+\tau} \kappa) \right) \delta_{t+s} \right] \quad (3)$$

To simplify this even further, we abstract from congestion effects, i.e. if $\delta_{t+\tau} = \delta$. In this case, $V_{i,t+1}$ is well approximated by a constant,

$$V_{i,t+1} \approx \frac{-\beta \delta \kappa u_d}{(1-\beta)(1-\beta((1-\rho-\delta\kappa)))} \equiv V_i^*$$

Under this approximation, the marginal utility of remaining susceptible becomes

$$V_{s,t+1} \approx \left[\gamma e_{t+1} \frac{I_{t+1}}{N} + \sum_{s=1}^{\infty} \beta^s \left(\prod_{\tau=0}^{s-1} \left(1 - \gamma e_{t+1+\tau} \frac{I_{t+1+\tau}}{N} \right) \right) \gamma e_{t+1+s} \frac{I_{t+1+s}}{N} \right] \beta V_i$$

Note that the term inside square brackets becomes larger when the risk of future infection, $\gamma e_{t+1+s} \frac{I_{t+1+s}}{N}$, goes up. Since V_i is constant (and negative), this implies that the risk of future infection *reduces* $V_{s,t+1} - V_{i,t+1}$, i.e. the private value of avoiding infection today is decreasing in the risk of future infection. In other words, anticipation of larger infection risks in the future creates perverse incentives for private agents. Why bother mitigating today if we are likely to become infected in the future? We term this the **fatalism effect**.

What happens with congestion effects, i.e. $\delta'(\kappa I_t) > 0$? Now, higher infections in the future push future mortality risk higher, which increases the disutility of getting infected today. However, the congestion effect has an even greater effect on the value of remaining susceptible. To put it differently, avoiding infection today is much less attractive if the healthcare system is expected to be congested in the future. This force, which we term the **front running effect**, also creates perverse incentives for

mitigation strategies in a laissez-faire equilibrium.

4 Planner's Problem

We normalize $N = 1$ for simplicity. The planner solves

$$\max U = \sum_{t=0}^{\infty} \beta^t u(C_t, L_t; I_t, D_t)$$

subject to

$$u(C_t, L_t; I_t, D_t) = (1 - D_t) \log(C_t) - (1 - D_t - \kappa I_t) \frac{L_t^{1+\eta}}{1+\eta} - u_\kappa \kappa I_t - u_d D_t$$

and

$$(1 - D_t) C_t = (1 - D_t - \kappa I_t) \left(L_t - \frac{\chi(\bar{M}_t)}{2} (M_t)^2 \right).$$

The first order conditions for consumption and labor are then (highlighted in red are the differences with the decentralized equilibrium)

$$\begin{aligned} C_t : C_t^{-1} &= \lambda_t + \textcolor{red}{2} \lambda_{e,t} e^c C_t \\ L_t : L_t^\eta &= \lambda_t - \textcolor{red}{2} \lambda_{e,t} e^l (1 - M_t)^2 L_t \\ M_t : \lambda_t \chi_t M_t &= \frac{\beta V_{\bar{M},t+1}}{1 - D_t - \kappa I_t} + \textcolor{red}{2} \lambda_{e,t} e^l (1 - M_t) L_t^2 \end{aligned}$$

The marginal utilities of the planner with respect to exposure are twice as high as those of the private sector because of the contagion externalities: private agents only care about how their behavior affect their own infection risk. They do not care about how their behavior affects the infection risk of others.

The envelope condition that changes (relative to the equilibrium) is the one with respect to the number of infected people:

$$\begin{aligned} V_{I,t} &= \kappa \frac{L_t^{1+\eta}}{1+\eta} - \kappa u_\kappa - \kappa \lambda_t (L_t - \chi_t M_t^2) + \lambda_{e,t} \kappa e^l (1 - M_t)^2 L_t^2 - (1 - \rho) \lambda_{i,t} \\ &\quad - \gamma e_t S_t \lambda_{i,t} - (\delta_t \kappa + \textcolor{red}{\delta'_t \kappa^2 I_t}) (\lambda_{d,t} - \lambda_{i,t}). \end{aligned}$$

4.1 Social Incentives to Mitigate

Let us now consider the planner's incentives to undertake mitigation today, summarized by the multiplier $\lambda_{e,t}$. As in equilibrium, this is the product of the current infection risk from exposure and the value of avoiding infection today.

$$\lambda_{e,t} = \beta (V_{S,t+1} - V_{I,t+1}) \gamma \frac{I_t}{N} S_t$$

As before, we focus on the empirically relevant region where the cost of death parameter, u_d is large. Then, the planner's marginal utility of today infection is

$$\begin{aligned} V_{I,t+1} \approx & -\frac{\beta u_d}{1-\beta} [\delta_{t+1} + \delta'_{t+1} \kappa I_{t+1} \\ & + \sum_{s=1}^{\infty} \beta^s \left(\prod_{\tau=0}^{s-1} (1 - \rho + \gamma e_{t+1+\tau} S_{t+1+\tau} - \delta_{t+1+\tau} - \delta'_{t+1+\tau} \kappa I_{t+\tau}) \right) (\delta_{t+1+s} + \delta'_{t+1+s} \kappa I_{t+1+s}) \Big]. \end{aligned} \quad (4)$$

Comparing this to the private value of infection in 3 reveals the effect of the externalities. The $\gamma e_{t+1+\tau} S_{t+1+\tau}$ term reflects the risk of a newly infected person spreading the disease to others in the future (intuitively, from the planner's perspective, the effective recovery rate is not ρ , but $\rho - \gamma e_t S_t$). This risk (increasing in the number of susceptible agents) makes the planner dislike current infection more than private agents do. The congestion externality $\delta'_{t+1+s} \kappa I_{t+1+s}$ also pushes up the social cost of current infections relative to the private costs. These two forces more than offset the fatalism and front running effects, causing the planner to mitigate aggressively when she perceives high risk of infections in the future.

The tension between the fatalism and precautionary effects might explain the Florida spring break controversy. As news about high future $I_{t+\tau}$, private citizens chose to enjoy their spring break, arguing that if they were going to catch the virus, now would be as good a time as later, while public officials worried about hospitals being overwhelmed.

5 Calibration

The lack of reliable data to calibrate the contagion model creates a serious challenge and an important limitation. Atkeson (2020a) discusses these difficulties. We calibrate our model at the weekly frequency.

Contagion The SIR block of the model is parameterized as follows. The recovery parameter is set to $\rho = 0.35$. The fraction of infected people who are sick is $\kappa = 0.15$. We normalize $\bar{e} + e^c + e^l = 1$. In our baseline calibration, we set the exposure loading parameters $e^c = e^l = \frac{1}{3}$ which is consistent with the estimation in Ferguson (2020). These parameters imply $e = 1$ at the pre-pandemic levels of consumption and labor (the calibration of production and utility parameters will be described later). The parameter γ is then chosen to target the basic reproduction number (i.e. the average number of people infected by a single infected individual) of $\mathcal{R} = 2$, yielding an estimated value of $\gamma = 0.7$. Finally, to parameterize the fatality rate and the congestion effects, we adopt the following functional form for $\delta(\cdot)$:

$$\delta(\kappa I_t) = \bar{\delta} + \exp(\phi I_t) - 1$$

where the parameter ϕ indexes the strength of the congestion externality. We set $\bar{\delta}$ and ϕ to match two targets for the case fatality rate: a baseline value (i.e. the fraction of infected people who die even in the absence of congestion) of 1% and an ‘extreme’ value (the fraction of people who die $\kappa I = 0.15(0.2)$, i.e. 3% of the population requires medical attention) of 5%. This procedure yields $\bar{\delta} = 0.023$ and $\phi = 3.15$.

Preferences and technology The utility parameter u_d is set to a baseline value of 2. This implies a flow disutility from death that is roughly 7 times per capita income. Such large non-monetary costs associated with loss of life are consistent with estimates in the literature and with values used by government entities like the EPA. For example, Greenstone and Nigam (2020) use an estimated value of a statistical life of \$11.5 million (in 2020 dollars) to the household from death. Assuming a rate of return of 5%, this translates into an annual flow value of \$575,000, or roughly 10 times per capita GDP. The flow disutility from sickness u_s is set to equal one-fourth of u_d , i.e. a value of 0.5. To calibrate the working-from-home technology, we make use of preliminary estimates for the effect of the lockdown in February and March on the Chinese economy. Recent data suggest a 25% fall in Industrial Value Added during the month of February. Assuming near complete lockdown, i.e. $m = 1$, this gives us an upper bound for $\bar{\chi}$, which indexes the initial productivity loss, of 0.5. Finally, for the long-run loss parameter, Δ_χ , we rely on Dingel and Neiman (2020), who estimate that roughly one-third of the jobs

in the US can be done from home. Guided by this estimate, we set $\Delta_\chi = 0.34$. We assume that the cumulation of learning depreciates very slowly $\bar{m}_{t+1} = 0.99\bar{m}_t + m_t$.

Initial Conditions, Vaccine, and Robustness A time period is interpreted as a week. The discount factor β captures both time discounting and the discovery of a cure/vaccine. We assume for simplicity that a cure and a vaccine arrive randomly together with a constant arrival rate. This is then exactly equivalent to adjusting β . We take a relatively pessimistic case as our baseline, where the combined effect of time discounting and the vaccine is to yield an annual β of 0.8 and a weekly beta of $\beta = (0.8)^{\frac{1}{52}} = 0.9957$.

The initial infection rate (the fraction of the infected people at time 0) is $i_0 = 1\%$ in our baseline. We run two robustness experiments, one with $i_0 = 0.1\%$, one with $\delta(\kappa I_t)/10$ to capture the uncertainty described in Atkeson (2020a).

6 Quantitative Results

Our benchmark exercise uses a large initial infection rate of $i_0 = 1\%$ because it makes the figures easier to read, but it is important to keep in mind that this is a large shock. It seems likely that agents and policy makers become aware of the epidemic much earlier so we also report simulations starting at $i_0 = 0.1\%$.

Private Response The figures show the results of simulations. We start with the decentralized solution. Figure 1 and Figure 2 show the behavior of the contagion and macro variables in the decentralized equilibrium, under 3 different assumptions about exposure and mitigation strategies. The blue line solid shows a situation where infection rates are exogenous, i.e. do not vary with the level of economic activity. Since infection is assumed to be exogenous, agents do not engage in mitigation, i.e., they ignore the pandemic. In fact, labor input rises (the solid line, top left panel in Figure 2), while per-capita consumption falls by about 2.5% (the dashed line), as able-bodied workers work harder to compensate for the workers who are sick. This is of course not a realistic assumption, but it serves as a useful benchmark for the worst case scenario. In this scenario, eventually about 80% of the population is infected and about 2.5% of the population succumbs to the virus (bottom, left panel in Figure 1). The case mortality rate peaks at 4% roughly 15 weeks after the initial infection because, at the peak, about 15% of the population is infected and the healthcare system is overwhelmed.

The red line describes the case where exposure is endogenous (i.e. varies with consumption and labor) but there is effectively no work-from-home (WFH) technology, i.e, the only ways for the household to reduce exposure are to cut back on its consumption and labor supply. As we would expect, this leads to a sharp reduction in economic activity (top, left panel in Figure 2) by about 10%. Importantly, however, the reduction is gradual, tracking the overall infection rate (it takes almost 17 weeks for consumption and labor to hit their trough). Intuitively, when the fraction of infected people is low (as is the case in the early stages), a reduction in exposure has a small effect on future infection risk, relative to the resulting fall in consumption. And since each household does not internalize the effect it has on the future infection rate, it has little incentive to indulge in costly mitigation early on. This dynamic is reflected in the hump-shaped pattern in λ_e (the bottom, middle panel in Figure 2). As we will see, this is drastically different in the planner’s problem. The mitigation behavior does lower the cumulative infection and death rates (relative to the exogenous infection risk) down do about 2%.

Finally, the yellow line shows the effect of the WFH technology. This allows the household to reduce exposure without sacrificing consumption – now, the peak loss in consumption is 7.5% (compared to 10% without the WFH option), even as the exposure falls by more (0.8 compared to 0.92, bottom left panel in Figure 2). The fraction of time spent working from home (top, left panel in Figure 2) is hump-shaped, peaking at 20% at roughly the same time as the fraction of infected rate. This additional flexibility also lowers cumulative fatalities to 1.75%. However, the timing of mitigation strategies is mostly unchanged – households in the decentralized economy do not find it optimal to front-load their mitigation efforts.

Optimal Response We now turn to the planner’s solution, depicted in Figures 3 and 4. As before, the blue, red and yellow lines show the cases of exogenous infection, mitigation without WFH and mitigation with WFH respectively. As the yellow and red curves in Figure 3 clearly show, the planner finds it optimal to “flatten the curve” rather dramatically. The peak infection and mortality rates are only slightly higher than their initial levels and well below the decentralized equilibrium levels, as are cumulative fatalities (approximately 0.6%, compared to 1.75% in the decentralized equilibrium with WFH). To achieve this, the planner has to reduce exposure drastically by more than 40% (recall that, in the decentralized equilibrium, exposure bottomed out at 0.8), keeping the basic reproduction number \mathcal{R} from rising much above 1. Of course, this pushes the economy into a deep recession with consumption falling by as much as 40% in the case without WFH (top left panel in Figure 4). More

interestingly, perhaps, the planner chooses to step on the brakes almost immediately, rather wait for infection rates to rise. In fact, the shadow value of exposure (bottom middle panel in Figure 4) spikes upon impact and then slowly decays over time, as the number of susceptible people declines.

The availability of the WFH technology ameliorates the economic impact of the planner’s suppression strategy, but does not significantly alter the contagion dynamics. Intuitively, fatalities are so costly that the planner aggressively suppresses the infection even in the absence of the WFH option: access to WFH simply allows the planner to achieve the same exposure outcomes at a lower cost. Notice also that the planner’s incentives to use the WFH technology are much stronger than that of private agents in a decentralized equilibrium: partly because her shadow value of exposure is higher, but also because she attaches a greater value to the future benefit from learning-by-doing. As a result, m rises to as high as 0.5 relatively quickly. In other words, as long as a vaccine is not available, the planner requires 40% of the population to work from home.

Smaller initial shock What is the value of an early warning? Suppose agents become aware of the disease at $i_0 = 0.1\%$. The private sector response, because it still follows the infection curve, barely changes. Private agents do not have the proper incentives to use the early warning. The planner, on the other hand, continues to front-load her effort and achieves a much better outcome in that case. The cumulative fatality rate is only 0.15%.

Lower fatality As we discussed earlier, there is little agreement in the literature on the mortality of Covid-19, with Atkeson (2020a) suggesting that the range of plausible values could be as wide as 0.1% to 1%. To explore the effect of this uncertainty, we repeat our analysis of the planner’s problem with mortality rates that are a tenth of our baseline values, i.e. probability of death (conditional on being) is now given by $\frac{1}{10}\delta(\kappa I_t)$. The results are presented in Figures 9 and 10. As expected, the overall number of fatalities are much lower, but the overall profile of the planner’s solution is unchanged: she still finds it optimal to front load her intervention. The economic cost, however, is much lower and the lockdown finishes earlier to let the economy recover.

Congestion externality Next, to isolate the role of the congestion externality, we repeat our analysis of the planner’s solution with the parameter ϕ set to 0. The results, shown in Figures 9 and 10, display a similar pattern as Figures 3 and 4, but with a less severe contraction. Intuitively, a healthcare system with sufficient slack capacity allows the planner to achieve similar outcomes in terms of fatali-

ties with modestly higher infection rates or equivalently, a higher level of exposure through economic activity (top, left panel in Figure 10). Accordingly, the recession is slightly less deep and much less persistent than in the baseline.

Fatalism and Front-running Figure 11 illustrates the perverse effects of fatalism and front-running on private incentives. In equilibrium (the two left panels), the onset of the pandemic pushes reduces the value of avoiding infection ($V_{s,t} - V_{i,t}$). This is particularly striking with congestion in the healthcare system (the bottom left panel), where $V_{s,t}$ is lower than $V_{i,t}$ early on. In other words, private agents might prefer to get infected today rather than remain vulnerable to infection at a future date. The corresponding plots for the planner (the two panels on the right) paint a very different picture – $V_{I,t}$ drops dramatically upon impact, reflecting the infection and congestion externalities imposed by current infections. The value of remaining susceptible drops too, but much more modestly. This explains the planner’s aggressive suppression of the disease right from the start.

Altruistic households Finally, we consider the possibility that households partly internalize their effects on the aggregate outcome. Formally, we assume that households incorporate externalities with a weight τ while making their consumption, labor and mitigation decisions. The *laissez-faire* equilibrium and the planner’s optimum are obtained as special cases by setting $\tau = 0$ and $\tau = 1$ respectively. If $\tau = 0.33$, the death rate is 1.4% (only slightly lower than the 1.7% under $\tau = 0$). If τ is assumed to be 0.5, i.e. households place equal weight on private and social incentives, the cumulative death rate falls to 1.25% (roughly a third of the gap between the *laissez-faire* equilibrium and the planner’s optimum). These results show that for private decisions to be close to that of the planner requires a very large degree of altruism.

7 Conclusions

We propose an extension of the neoclassical model to include contagion dynamics, to study and quantify the tradeoffs of policies that can mitigate the Covid-19 pandemic. Our model reveals two key insights. First, that relative to the incentives of private agents, a planner wishes to significantly front-load mitigation strategies. Second, in our calibrated model, the prospect of mitigation together with the possibility of agents working from home gives quantitatively meaningful reductions in the spread of a disease and the economic costs.

As we write the first draft of this paper, there is much uncertainty about the parameters of the disease, and yet decisions must be made. Some of our results speak directly to this dilemma. Atkeson (2020a) points out that, when one does not know the initial number of active cases, it is difficult “to distinguish whether the disease is deadly (1% fatality rate) or milder (0.1% fatality rate).” In our simulations we have considered a deadly disease with a low initial infection rate of $i_0 = 0.1\%$, and a milder disease with a high initial infection rate of $i_0 = 1\%$. Interestingly, in both cases, the planner should implement immediately a strong suppression policy. The main difference is that in the mild case it is optimal to release the lockdown sooner. Thus, faced with large uncertainty about the fatality rate, the planner could impose a strong suppression policy and wait for more data. Assuming data available 20 weeks after the outbreak allows her to precisely estimate the fatality rates, this strategy would allow the planner to implement the optimal response, despite the initial uncertainty about a key parameter.

Extensions of our baseline setup would yield a fruitful laboratory to study other policy-relevant questions. For instance, the production-side of the economy could be enriched to multiple and heterogeneous sectors, to provide additional margins to mitigate exposure and economic hardship. A richer model will also help to give sharper quantitative estimates on the costs of mitigation strategies.

References

- Alvarez, Fernando, David Argente, and Francesco Lippi**, “A Simple Planning Problem for COVID-19 Lockdown,” March 2020. Working Paper.
- Atkeson, Andrew**, “How Deadly is COVID-19? Understanding the Difficulties with Estimation,” March 2020. NBER Working Paper No. 26867.
- , “What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios,” March 2020. NBER Working Paper No. 26867.
- Baker, Scott R., R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis**, “How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic,” *University of Chicago, Working Paper*, 2020.

- Barro, Robert J, Jose Ursua, and Joanna Weng**, “The coronavirus and the Great Influenza epidemic: Lessons from the “Spanish Flu” for the coronavirus’ potential effects on mortality and economic activity,” March 2020. AEI Economics Working Paper 2020-02.
- Berger, David, Kyle Herkenhoff, and Simon Mongey**, “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” March 2020. Working Paper.
- Correia, Sergio, Stephan Luck, and Emil Verner**, “Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu,” March 2020. Working Paper.
- Diekmann, Odo and J. Heesterbeek**, *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*, John Wiley, 2000.
- Dingel, Jonathan and Brent Neiman**, “How Many Jobs Can be Done at Home?,” March 2020. Working Paper.
- Eichenbaum, Martin, Sergio Rebelo, and Mathias Trabandt**, “The Macroeconomics of Epidemics,” March 2020. Working Paper NWU.
- Ferguson, Neil**, “Impact of non-pharmaceutical interventions (NPIs) to reduce COVID- 19 mortality and healthcare demand,” March 2020. Imperial College Report.
- Greenstone, Michael and Vishan Nigam**, “Does Social Distancing Matter?,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020, (2020-26).
- Harko, Tiberiu, Francisco Lobo, and M. K. Mak**, “Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates,” *Applied Mathematics and Computation*, 2014, (236), 184–194.
- Lucas, Robert E. Jr. and Nancy L. Stokey**, “Money and Interest in a Cash-in-Advance Economy,” *Econometrica*, 1987, 55 (3), 491–513.

Figure 1: Decentralized Equilibrium: Contagion Dynamics

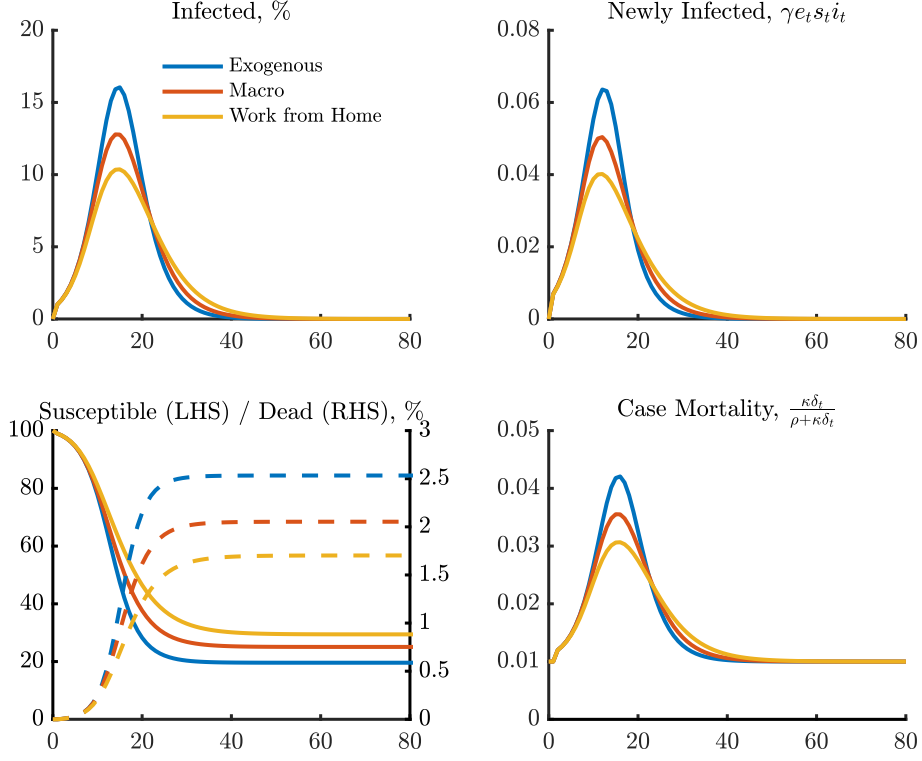


Figure 2: Decentralized Equilibrium: Aggregates

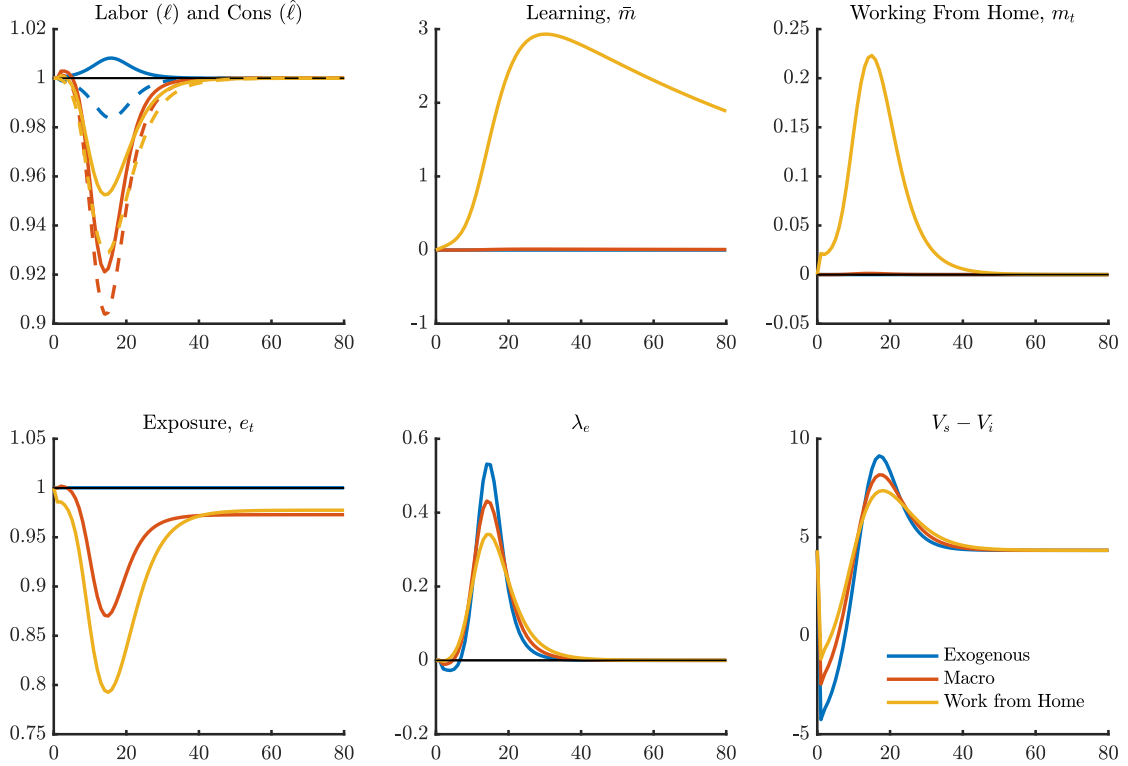


Figure 3: Planner Solution: Contagion Dynamics

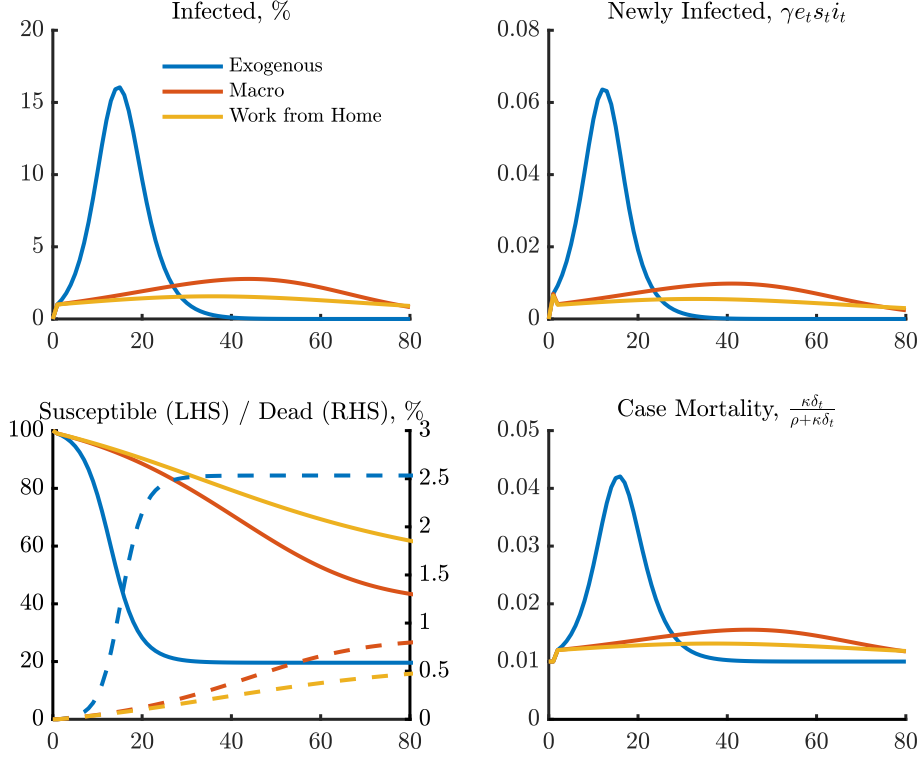


Figure 4: Planner Solution: Aggregates

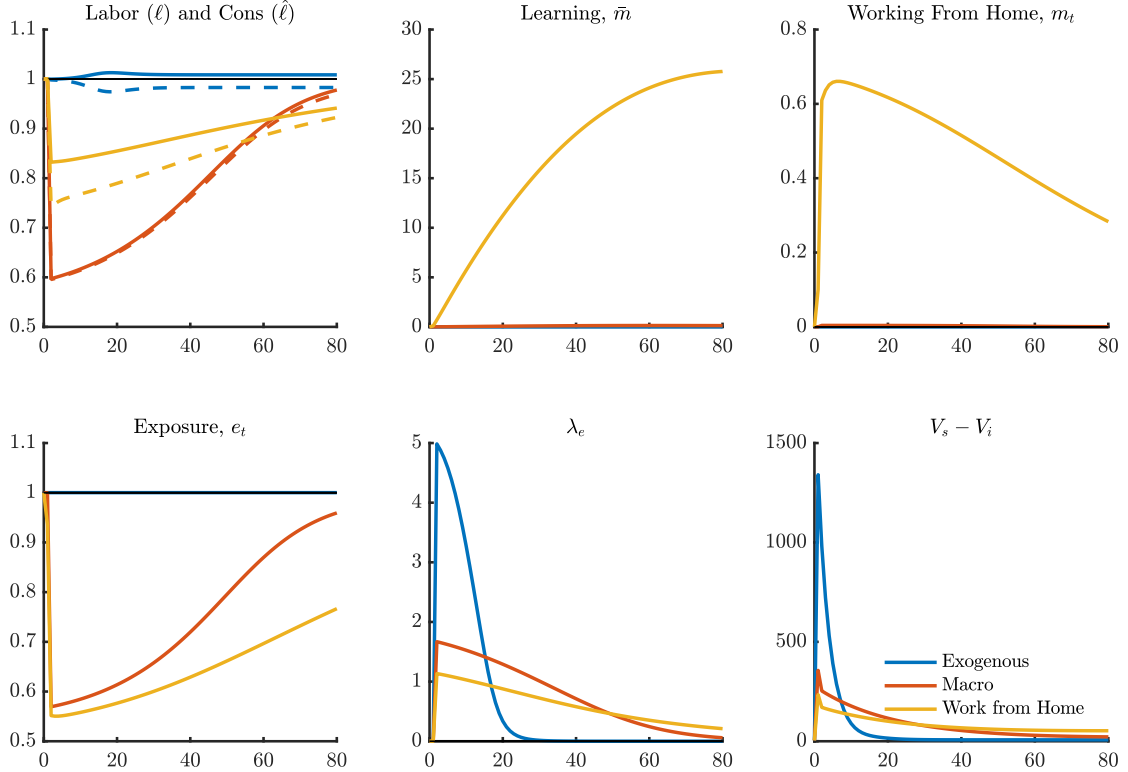


Figure 5: Planner Solution: Contagion Dynamics with Smaller Infection Shock

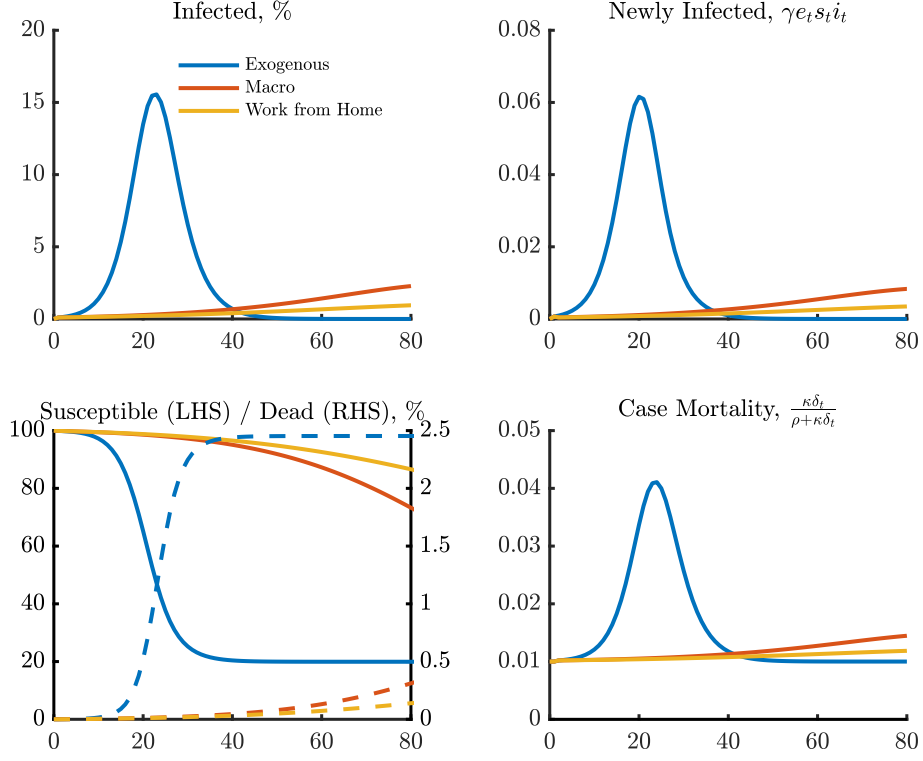


Figure 6: Planner Solution: Aggregates with Smaller Infection Shock

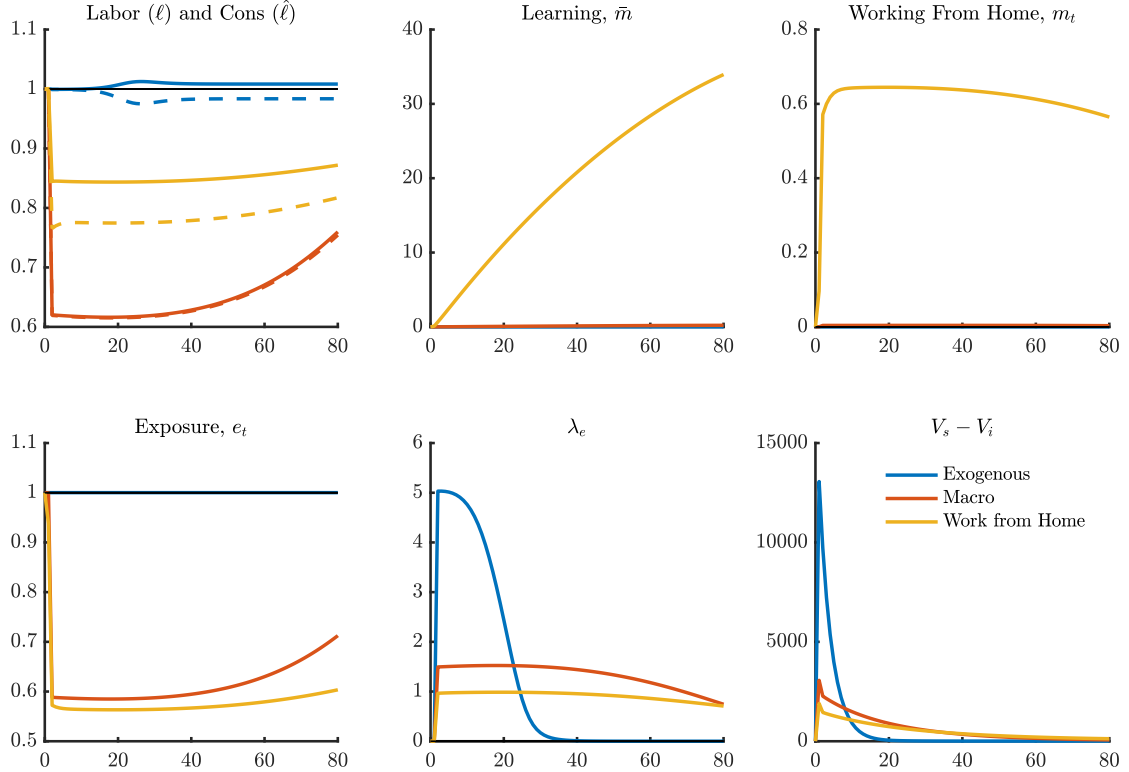


Figure 7: Planner Solution: Contagion Dynamics with Reduced δ_t

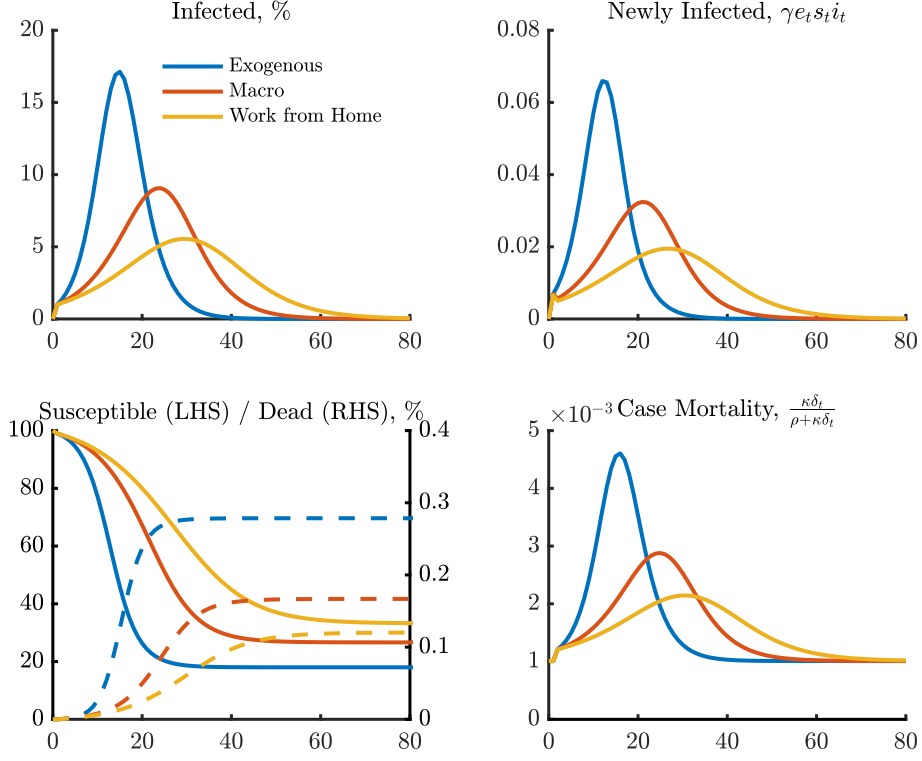


Figure 8: Planner Solution: Aggregates with Reduced δ_t

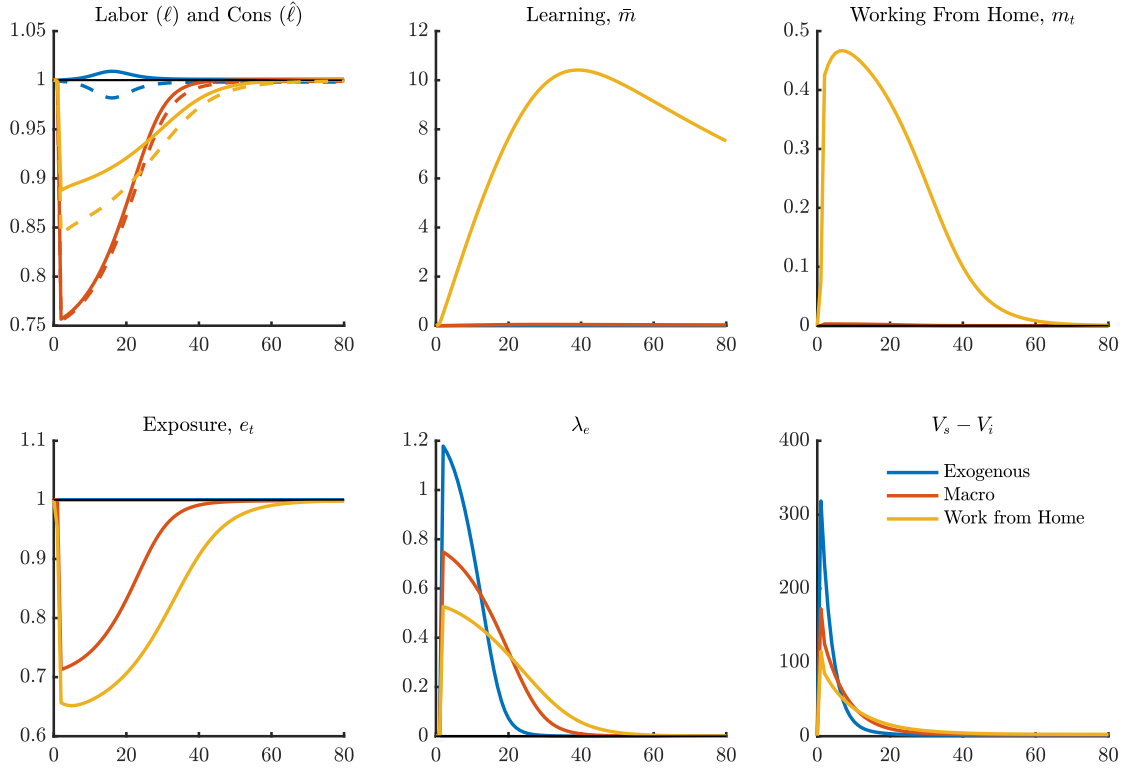


Figure 9: Planner Solution: Contagion Dynamics with $\phi = 0$

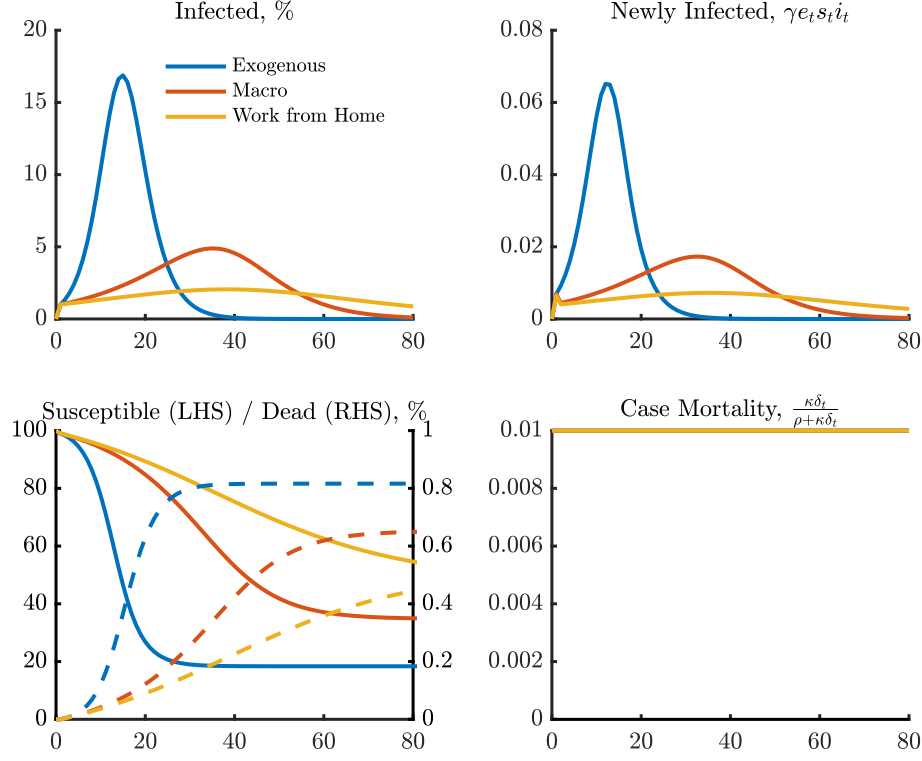


Figure 10: Planner Solution: Aggregates with $\phi = 0$

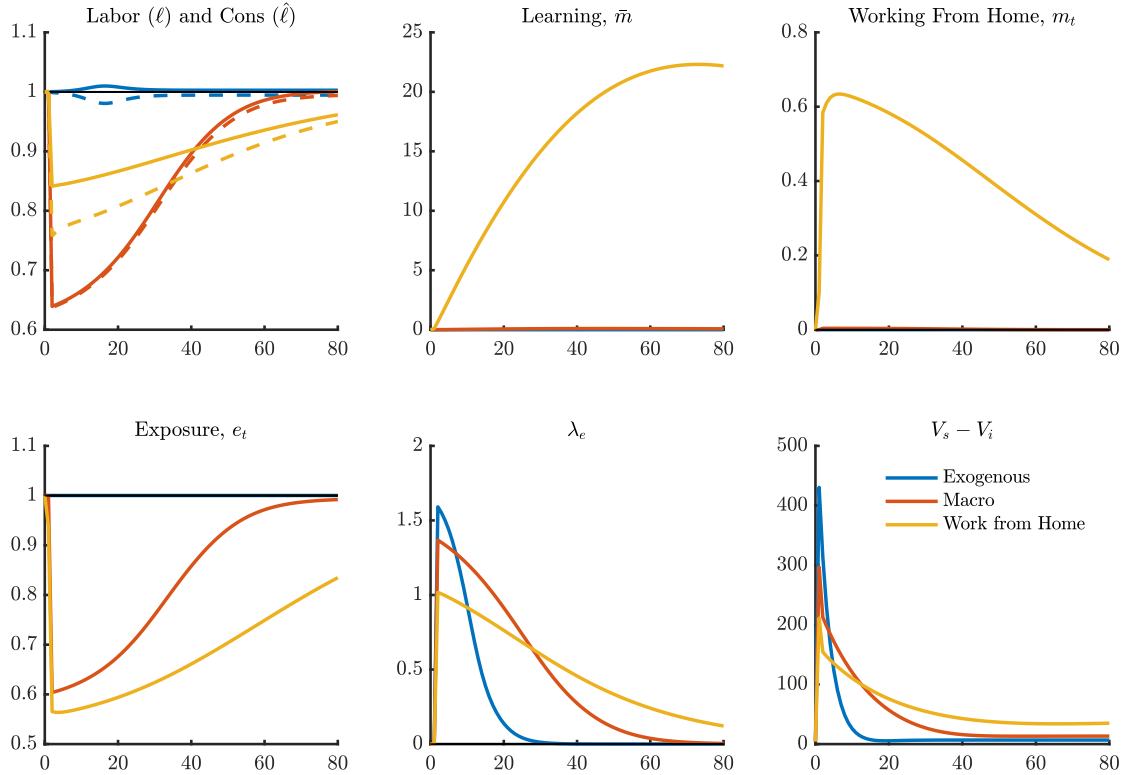
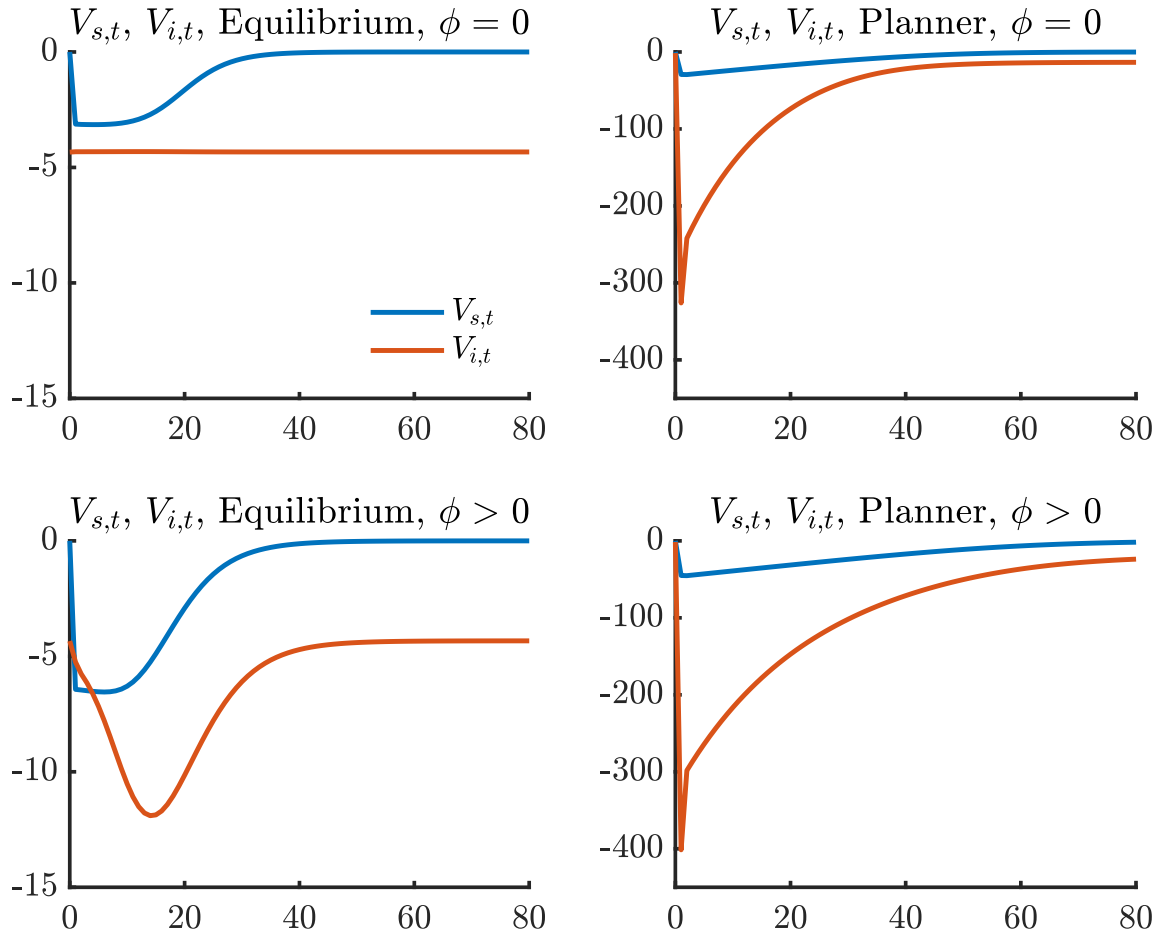


Figure 11: Fatalism and Front-running



Appendix

A Planner's Problem

The planner understands that the dynamics follow

$$\begin{aligned} S_{t+1} &= S_t - \gamma e_t I_t S_t \\ I_{t+1} &= \gamma e_t I_t S_t + (1 - \rho) I_t - \delta (\kappa I_t) \kappa I_t \\ D_{t+1} &= D_t + \delta (\kappa I_t) \kappa I_t \\ R_{t+1} &= R_t + \rho I_t \end{aligned}$$

We normalize $N = 1$ for simplicity.

$$\max U = \sum_{t=0}^{\infty} \beta^t u(C_t, L_t; I_t, D_t)$$

subject to

$$u(C_t, L_t; I_t, D_t) = (1 - D_t) \log(C_t) - (1 - D_t - \kappa I_t) \frac{L_t^{1+\eta}}{1+\eta} - u_\kappa \kappa I_t - u_d D_t$$

and

$$(1 - D_t) C_t = (1 - D_t - \kappa I_t) \left(L_t - \frac{\chi(\bar{M}_t)}{2} (M_t)^2 \right)$$

The Lagrangian is

$$\begin{aligned} V_t(I_t, S_t, D_t, \bar{M}_t) = & u_t + \beta V_{t+1} + \lambda_t \left((1 - D_t - \kappa I_t) \left(L_t - \chi(\bar{M}_t) (M_t)^2 \right) - (1 - D_t) C_t \right) \\ & + \lambda_{e,t} \left(e_t - \bar{e} - (1 - D_t) e^c C_t^2 - (1 - D_t - \kappa I_t) e^l (1 - M_t)^2 L_t^2 \right) \\ & + \lambda_{i,t} (I_{t+1} - \gamma e_t I_t S_t - (1 - \rho) I_t + \delta (\kappa I_t) \kappa I_t) \\ & + \lambda_{s,t} (S_{t+1} - S_t + \gamma e_t I_t S_t) \\ & + \lambda_{d,t} (D_{t+1} - D_t - \delta (\kappa I_t) \kappa I_t) \end{aligned}$$

The first order conditions for consumption and labor are then (highlighted in red the difference with the decentralized equilibrium)

$$\begin{aligned}
C_t : C_t^{-1} &= \lambda_t + \textcolor{red}{2}\lambda_{e,t}e^c C_t \\
L_t : L_t^\eta &= \lambda_t - \textcolor{red}{2}\lambda_{e,t}e^l (1 - M_t)^2 L_t \\
M_t : \lambda_t \chi_t M_t &= \frac{\beta V_{\bar{M},t+1}}{1 - D_t - \kappa I_t} + \textcolor{red}{2}\lambda_{e,t}e^l (1 - M_t) L_t^2
\end{aligned}$$

The marginal utilities of the planner with respect to exposure are twice as high as those of the private sector because of the contagion externalities: private agents only care about how their behavior affect their own infection risk. They do not care about how their behavior affects the infection risk of others.

The remaining first order conditions are the same as those of the private sector

$$\begin{aligned}
e_t : \lambda_{e,t} &= (\lambda_{i,t} - \lambda_{s,t}) \gamma I_t S_t \\
I_{t+1} : \lambda_{i,t} &= -\beta V_{I,t+1} \\
S_{t+1} : \lambda_{s,t} &= -\beta V_{S,t+1} \\
D_{t+1} : \lambda_{d,t} &= -\beta V_{D,t+1}
\end{aligned}$$

The envelope conditions are

$$\begin{aligned}
V_{I,t} &= \kappa \frac{L_t^{1+\eta}}{1+\eta} - \kappa u_\kappa - \kappa \lambda_t \left(L_t - \frac{\chi_t}{2} M_t^2 \right) + \lambda_{e,t} \kappa e^l (1 - M_t)^2 L_t^2 - (1 - \rho) \lambda_{i,t} \\
&\quad - \gamma e_t S_t \lambda_{i,t} - (\delta_t \kappa + \textcolor{red}{\delta'_t \kappa^2 I_t}) (\lambda_{d,t} - \lambda_{i,t}) \\
V_{S,t} &= -\lambda_{s,t} - \gamma e_t I_t (\lambda_{i,t} - \lambda_{s,t}) \\
V_{D,t} &= \frac{L_t^{1+\eta}}{1+\eta} - \log(C_t) - u_d - \lambda_t \left(L_t - \frac{\chi_t}{2} (M_t)^2 - C_t \right) + \lambda_{e,t} \left(e^c C_t^2 + e^l (1 - M_t)^2 L_t^2 \right) - \lambda_{D,t} \\
V_{\bar{M},t} &= \beta V_{\bar{M},t+1} + \lambda_t (1 - D_t - \kappa I_t) \frac{\bar{\chi}}{2} \Delta_\chi e^{\bar{M}_t} (M_t)^2
\end{aligned}$$

B Properties of Contagion Dynamics

B.1 Definitions

We start with the most basic concept in epidemiology, the basic reproduction number, which we denote by \mathcal{R} because the usual notation “ \mathcal{R}_{not} ” is terribly confusing. \mathcal{R} is the expected number of cases

directly generated by one case when everyone else is susceptible. The most basic model is to assume that when someone is infected there are three stages

1. a latency period T_1 when the individual is not yet infectious
2. infectious period $T_2 - T_1$
3. recovered period after T_2 when the individual is not infectious anymore

If the contact rate (exposure) is e and the probability of infection conditional on contact is γ , the expected number of secondary cases per primary case in a fully susceptible population is therefore

$$\mathcal{R} = \gamma e (T_2 - T_1)$$

In our notations, e is the number of people that one individual meets per unit of time and γ is the probability of transmitting the disease conditional on a meeting between one infectious and one susceptible agent.

B.2 SIR Model

The *SIR* model builds on this idea. We define the length of one period so that $T_1 = 1$. If someone is infected in period t , then she will start spreading the disease in period $t + 1$. Let I_t be the number of infected individuals and S_t the number of susceptible individuals at the beginning of time t in a population of size N . Each infected agent meets e people. We assume that the meetings are random and that the population is always evenly mixed, therefore the probability of meeting a susceptible person is S_t/N . The number of meetings between infected and susceptible agents is therefore $e I_t S_t/N$ and the total number of new infections is $\gamma e I_t S_t/N$. In our macro model e is an endogenous variable but we take it as a constant for now. We assume that recovery follows a Poisson process with intensity ρ . The infection equation is then

$$I_{t+1} = \gamma e I_t \frac{S_t}{N} + (1 - \rho) I_t \quad (5)$$

Consider a population of size N (large) of initially susceptible individuals ($S_0 = N$). If one individual is infected, the total number of secondary infections from that individual is

$$\mathcal{R} = \sum_{\tau=0}^{\infty} \gamma e (1 - \rho)^{\tau} = \frac{\gamma e}{\rho} \quad (6)$$

Note that \mathcal{R} is a number, not a rate per unit of time. The model has a steady state at $I = 0$ and $S = 1$ but it is unstable in the sense that if one individual gets infected the system converges to a different steady state. In general we can write

$$\frac{I_{t+1}}{I_t} = 1 + \gamma e \frac{S_t}{N} - \rho$$

When $S/N \approx 1$, the number infected people evolves exponentially as $\frac{I_{t+1}}{I_t} \approx 1 + \gamma e - \rho$. If $R_0 < 1$ then a small infection disappears exponentially. If $R_0 > 1$ then there is an epidemic where I initially grows over time. The growth continues as long as $\gamma e \frac{S_t}{N} > \rho$. Eventually the number of susceptible people decreases and growth slows down or reverses, depending on how we close the model.

There are two ways to close the model. The simpler one, called the *SIS* model, assumes that recovered agents ρI go back to the pool of susceptible agents. This is the model used to study the common cold. In that case $N = S_t + I_t$ and the equation becomes

$$I_{t+1} = \gamma e I_t \frac{N - I_t}{N} + (1 - \rho) I_t$$

and the steady state infection rate is $\frac{I}{N} = \max\left(0; 1 - \frac{\rho}{\gamma e}\right)$.

The other way is to introduce a population of recovered agents R who are not susceptible anymore. This model – called *SIR* – is used for flu epidemics, among others. In the simple model, R is an absorbing state. The system becomes

$$S_{t+1} - S_t = -\gamma e I_t \frac{S_t}{N} \tag{7}$$

$$R_{t+1} - R_t = \rho I_t \tag{8}$$

and of course $N = S_t + I_t + R_t$. Note that S is (weakly) decreasing and R is (weakly) increasing, therefore their limits exist. Since $N = S_t + I_t + R_t$ so does the limit of I . For R and S to be constant it must be that I tends to zero. Therefore

$$\lim_{t \rightarrow \infty} I_t = 0$$

That is the first simple property of the solution. Second, since S is decreasing, I must be (at most) single-peaked. If I_0 is small and $R_0 > 1$ then I_t must grow, reach a maximum, and then decrease

towards zero. This is the typical shape of the curves found in the literature. Harko et al. (2014) provide an analytical solution to the differential equations (in continuous time).

Let us now study the long run behavior of S and R . Combining equations (7) and (8) we get

$$R_{t+1} - R_t = -N \frac{\rho}{\gamma e} \frac{S_{t+1} - S_t}{S_t}$$

This equation is simpler to write in continuous time

$$\frac{\dot{S}}{S} = -\frac{\gamma e}{\rho} \frac{\dot{R}}{N}$$

and to integrate the solution:

$$\log \left(\frac{S_t}{S_0} \right) = -\frac{\gamma e}{\rho} \left(\frac{R_t - R_0}{N} \right)$$

This equation holds along any transition path without exogenous shocks. In the limit, since $S_\infty + R_\infty = 1$ we have the transcendental equation

$$\log \left(\frac{S_\infty}{S_0} \right) = -\frac{\gamma e}{\rho} \left(\frac{1 - S_\infty - R_0}{N} \right)$$

The long run steady state (S_∞, R_∞) depends on the initial conditions as well as the basic reproduction number \mathcal{R} . We can summarize our discussion in the following Lemma.

Lemma. *The SIR model is fully characterized by $\mathcal{R} = \frac{\gamma e}{\rho}$ and the initial conditions (S_0, R_0) . If $\mathcal{R} < 1$, infections die out without epidemic. If $\mathcal{R} > 1$, a small infection I_0 creates an epidemic: I_t rises, reaches a maximum in finite time before declining towards zero: $I_\infty = 0$. The long run limits S_∞ and R_∞ exist and satisfy $S_\infty + R_\infty = 1$ and $\log \left(\frac{S_\infty}{S_0} \right) = -\mathcal{R} \left(\frac{1 - S_\infty - R_0}{N} \right)$.*

The complete model takes into account that some individuals will die from the disease. We assume that a fraction κ of infected agents become (severely) sick and a fraction δ of the sick patients die. Hence we have another absorbing state, D . The system of equation of the *SIRD* model becomes

$$I_{t+1} = \gamma I_t \frac{S_t}{N} + (1 - \rho - \delta \kappa) I_t$$

$$S_{t+1} = S_t - \gamma I_t \frac{S_t}{N}$$

$$R_{t+1} = R_t + \rho I_t$$

$$D_{t+1} = D_t + \delta \kappa I_t$$

The number of sick people is κI_t and determines the pressure on the health care system. From the perspective of the epidemic we could aggregate D and R into one absorbing state: $\tilde{R} = D + R$ such that $\tilde{R}_{t+1} = \tilde{R}_t + (\rho + \delta\kappa) I_t$. The long run solution is

$$\log\left(\frac{S_\infty}{S_0}\right) = -\frac{\gamma e}{\rho + \delta\kappa} \left(\frac{1 - S_\infty - \tilde{R}_0}{N}\right)$$

and $\tilde{R}_\infty = 1 - S_\infty$ while $D_\infty = D_0 + \frac{\delta\kappa}{\delta\kappa + \rho} (\tilde{R}_\infty - \tilde{R}_0)$. From an economic and social perspective we need to keep track of D and R separately in any case.

B.3 SIR model with exogenous birth and death

The path dependence of the long run steady state is a somewhat artificial consequence of the lack of entry and exit. Suppose that ϵN people are both in state S each period, and also that there is a constant exogenous death rate ϵ . The system is

$$\begin{aligned} I_{t+1} &= \gamma e I_t \frac{S_t}{N} + (1 - \rho - \epsilon) I_t \\ S_{t+1} &= (1 - \epsilon) S_t - \gamma e I_t \frac{S_t}{N} + \epsilon N \\ R_{t+1} &= (1 - \epsilon) R_t + \rho I_t \end{aligned}$$

Note that population is constant: $S_{t+1} + I_{t+1} + R_{t+1} = N$. Now the steady state requires

$$\begin{aligned} \gamma e \frac{S}{N} I &= I(\epsilon + \rho) \\ \gamma e I \frac{S}{N} &= \epsilon(N - S) \\ \rho I &= \epsilon R \end{aligned}$$

Since $I > 0$ we can easily solve for the unique steady state

$$\begin{aligned} \frac{S}{N} &= \frac{\epsilon + \rho}{\gamma e} \\ \frac{I}{N} &= \epsilon \frac{1 - \frac{\epsilon + \rho}{\gamma e}}{\epsilon + \rho} \\ \frac{R}{N} &= \rho \frac{1 - \frac{\epsilon + \rho}{\gamma e}}{\epsilon + \rho} \end{aligned}$$

And now we can take the limit as $\epsilon \rightarrow 0$ to get $\frac{I}{N} = 0$, $\frac{S}{N} = \frac{\rho}{\gamma e} = \mathcal{R}^{-1}$ and $\frac{R}{N} = 1 - \frac{\rho}{\gamma e} = 1 - \mathcal{R}^{-1}$. Adding a small amount of exogenous birth and death renders the long run steady state independent of initial conditions.